



RESEARCH DEPARTMENT



REPORT

**DIGITAL VIDEO BIT RATE REDUCTION:
blanking elimination
for sub-Nyquist ($2f_{sc}$) PAL**

R.P. Marsden, B.A., M.B.K.S.

**DIGITAL VIDEO BIT RATE REDUCTION:
BLANKING ELIMINATION FOR SUB-NYQUIST ($2f_{sc}$) PAL**

R.P. Marsden, B.A., M.B.K.S.

Summary

At present the BBC transmits television signals between its United Kingdom premises using wide-band analogue circuits provided by British Telecom. As part of its modernisation plan British Telecom is converting its communications network to digital operation, and the analogue circuits will eventually disappear. Although video signals can easily be converted to digital form, the bit rate occupied by a conventional pulse-code modulation (p.c.m.) video signal is high. Economies in the use of circuits could be made by reducing this bit rate, although any reduction would have to be made without significantly affecting the television picture quality.

The Report discusses the elimination of blanking intervals from the digitised PAL signal, which is one of several techniques investigated as a means of reducing bit rates. Eliminating line and field blanking intervals reduces the bit rate by about 23%, and, associated with sub-Nyquist sampling and differential pulse code modulation (DPCM), would allow a broadcast quality television signal, with sound, to be transmitted at the third-order multiplex rate of 34.368 Mbit/s.

The equipment built for the study is described and its performance is discussed. Microprocessor control has been used, partly to allow changes to be made during development and partly because of the complexity of operation of the equipment. This complexity arises because tolerances on the PAL analogue signal that are acceptable in analogue processing and transmission must be carefully allowed for in digital processing.

It is concluded that blanking elimination, although possible, is unattractive for PAL-coded signals because of the complexity of the equipment. However, for signals that are more tightly defined, for example those complying with the CCIR recommendations for the digital coding of component video signals, blanking elimination would be a straightforward process that could be used with confidence to provide a valuable saving of bit rate with no impairment to picture quality.

Issued under the Authority of



Head of Research Department

**Research Department, Engineering Division,
BRITISH BROADCASTING CORPORATION**

June 1983
(EL-168)

This Report may not be reproduced in any form without the written permission of the British Broadcasting Corporation.

It uses SI units in accordance with B.S. document PD 5686.

DIGITAL VIDEO BIT RATE REDUCTION: BLANKING ELIMINATION FOR SUB-NYQUIST ($2f_{sc}$) PAL

Section	Title	Page
	Summary	Title Page
1.	Introduction	1
2.	Defining the blanking intervals	2
3.	Description of the equipment	4
	3.1. Coder	4
	3.1.1. Hardware	4
	3.1.2. The coder control program	6
	3.2. Decoder	9
	3.2.1. Hardware	9
	3.2.2. The decoder control program	11
4.	Performance	12
	4.1. Errors in input signal level	13
	4.2. Noise on analogue input	13
	4.3. Variations in line and field frequency	13
	4.4. Non-synchronous 'cuts'	13
	4.5. Propagation delay	14
5.	Conclusions	14
6.	References	14

© BBC 2006. All rights reserved. Except as provided below, no part of this document may be reproduced in any material form (including photocopying or storing it in any medium by electronic means) without the prior written permission of BBC Research & Development except in accordance with the provisions of the (UK) Copyright, Designs and Patents Act 1988.

The BBC grants permission to individuals and organisations to make copies of the entire document (including this copyright notice) for their own internal use. No copies of this document may be published, distributed or made available to third parties whether by paper, electronic or other means without the BBC's prior written permission. Where necessary, third parties should be directed to the relevant page on BBC's website at <http://www.bbc.co.uk/rd/pubs/> for a copy of this document.

DIGITAL VIDEO BIT RATE REDUCTION: BLANKING ELIMINATION FOR SUB-NYQUIST ($2f_{sc}$) PAL

R.P. Marsden, B.A., M.B.K.S.

1. Introduction

At present the BBC transmits television signals between its United Kingdom premises on wide-band analogue circuits provided by British Telecom. As part of its modernisation plan British Telecom is in the process of converting its communications network to digital operation, and these analogue circuits will eventually disappear.

For television pictures of broadcast quality the video signal must be quantised to at least 8-bit accuracy at a minimum sampling rate of about 12 MHz^1 . For PAL-coded signals a sampling rate of three times colour-subcarrier frequency is often used. With 8-bit representation this produces a bit rate of just over 106.4 Mbit/s . Adding allowances for sound, ancillary signals and error protection increases this to about 120 Mbit/s .

The British Telecom network will be based on a voice channel of 64 kbit/s , following the CCITT recommendation for a digital hierarchy, which is shown in Fig. 1. Access to the network will normally be at 2048 kbit/s , 8448 kbit/s , 34368 kbit/s and 139264 kbit/s ("140 Mbit/s") and possibly also at 68736 kbit/s . Without compression a digital PAL television signal would therefore take up almost all of a 140 Mbit/s trunk circuit.

With a view to achieving economy in the use of digital transmission an investigation was made into several methods of bit-rate reduction. Because the video signal is by far the largest part of the complete signal the work was concentrated on reducing the video bit rate.

Earlier work, using a relatively simple form of differential pulse-code modulation (DPCM)², had shown that the video bit-rate could be reduced

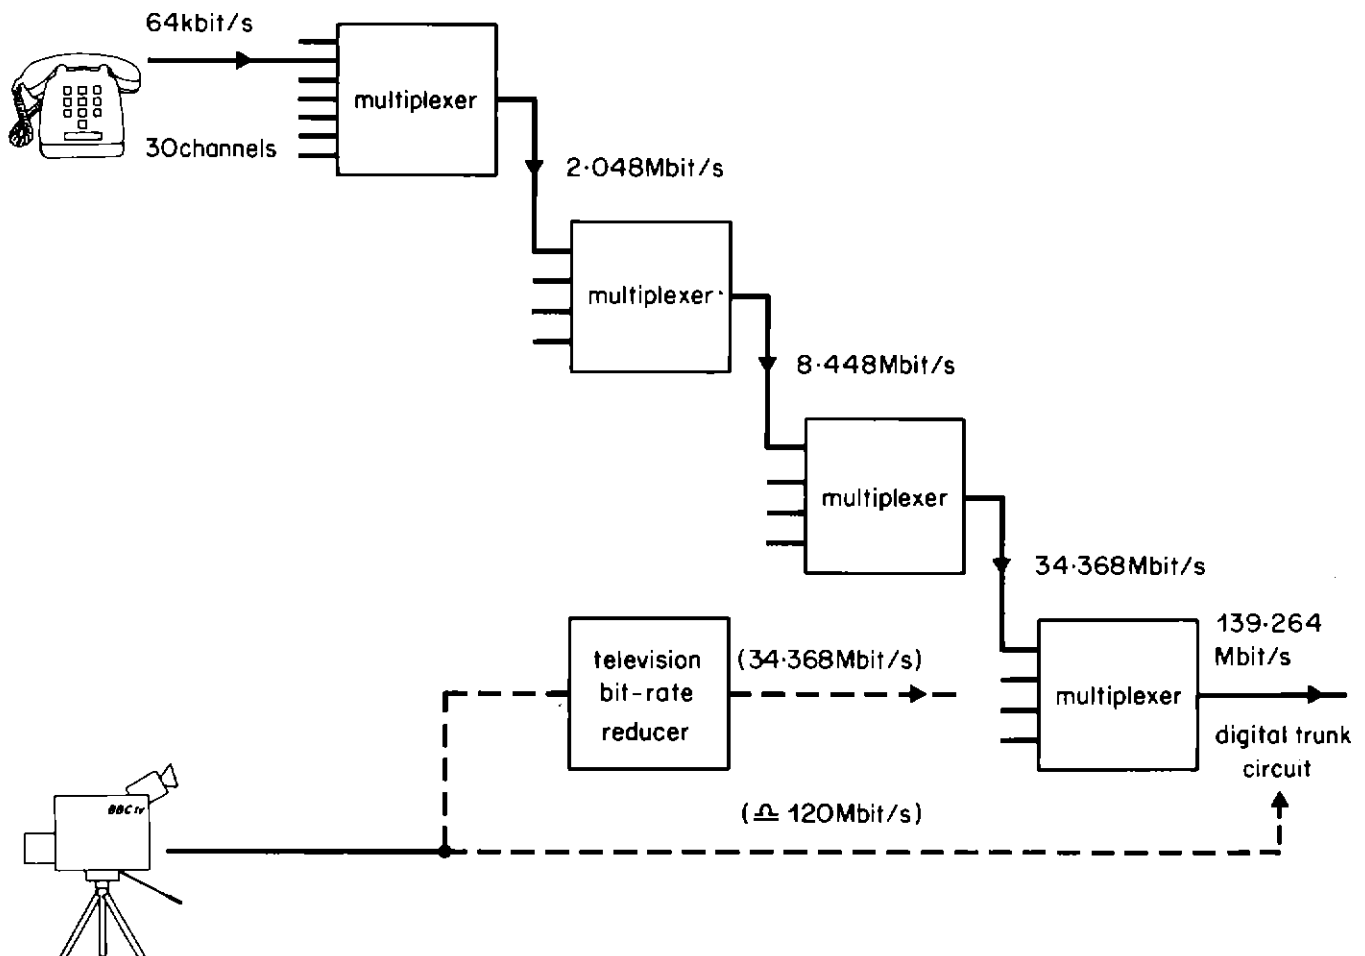


Fig. 1 – Multiplex hierarchy in a digital communications network.

by three eighths with an almost imperceptible effect on the picture quality. It was decided to develop this technique, and to investigate any others that might be applicable.

It was decided³ to attempt to encode a composite PAL video signal, with two sound channels, ICE* and TELETEXT within a total bit rate of 34368 kbit/s, the third-order multiplex rate. This would allow up to four television signals to be carried on one 140 Mbit/s circuit.

From preliminary work three techniques were chosen for further study:

- (a) sub-Nyquist sampling⁴, by which the sampling frequency would be reduced by 33% to twice colour-subcarrier frequency;
- (b) differential pulse-code modulation (DPCM)⁵, by which the number of bits needed to describe each sample would be reduced to $4\frac{1}{2}$ (actually by using 9 bits to represent two samples);
- (c) non-transmission of the line- and field-blanking intervals, which together make up almost one quarter of the video signal content.

A combination of these techniques would allow the video signal to be encoded in about 32 Mbit/s, including 2 Mbit/s for error protection.

This Report describes the third of these techniques, elimination of the blanking intervals, and the equipment that was designed and built for the investigation. The blanking intervals form about 25% of the video signal, and after allowing for tolerances on their position and for the transmission of some form of vestigial synchronising information a net saving of about 23% might be expected.

The main attraction of blanking-interval elimination is that it should not affect the quality of the picture. Samples from active parts of the picture are not changed in any way, and if the blanking intervals are correctly reconstituted the received signal should be indistinguishable from the original. In practice there are problems to be overcome. Some are fundamental, some practical and some arise from various aspects of BBC operating practice.

* 'ICE' is an acronym for 'Insertion Communication Equipment'. The waveform comprises a 126-bit binary pulse train, injected on lines 16 and 329. It is used for such purposes as identifying the signal and its source, and for transmitting monitoring data.

The colour subcarrier frequency, and hence the sampling frequency, is not an exact multiple of line frequency. The positions of the sampling points therefore vary from line to line.

Moreover, in source-synchronising operations the line and field frequencies will deviate from their nominal values while the frequency of colour subcarrier stays the same. This alters the number of samples taken in a field or line period. The blanking-elimination process must therefore detect these changes in frequency, and modify the pattern of samples that it is removing from the signal. Appropriate signals must then be sent to the receiving terminal so that the blanking intervals may be correctly replaced in all aspects. Thus, it is also necessary to measure the synchronising pulses and colour burst so that they can be reinserted with their original amplitudes. This is important as they may be used as references for automatic equalisation following subsequent transmission through an analogue circuit.

Because of these and other problems it was decided that a flexible design was required and since this would be difficult to achieve with conventional circuitry, the coder and decoder operations are controlled by fast microprocessors. This approach transferred much of the design and development work from the hardware itself to the writing of software and it also allowed modifications to be made in the light of experience without any physical rebuilding of the equipment.

2. Defining the blanking intervals

Although the nominal duration of the blanking intervals is well defined it is necessary to decide how much it is worthwhile or possible to remove.

Line blanking accounts for about 18% of the composite video waveform, field blanking for only a further 7%. The size of buffer store required to eliminate both field and line blanking is more than 100 times that required if only the line blanking intervals are removed. So it is reasonable to ask if it is worthwhile, in a 34 Mbit/s system, to remove the field blanking intervals.

The answer, basically, is yes. The target of 34 Mbit/s for a video signal and ancillaries is not easy to reach. With all blanking intervals removed 4.5 bits are available to code each sample, (assuming that twice subcarrier-frequency sampling is used). With only line blanking removed this falls to about 4.2 bits, which reduces the number of quantising levels from 22 to 17 for each sample.




Fig. 2 — The line blanking interval.

A more elaborate DPCM coder would thus be required, and the picture quality would almost certainly be worse. Alternatively, reducing the sampling frequency to maintain 4.5 bits per sample would lose the very considerable advantages of twice subcarrier-frequency sampling.

Although the larger buffer store would cost more, the cost, size and power consumption of storage continues to fall. The larger store does not incur the penalty that might first appear, particularly when the alternatives just described are considered. Also, the cost of the control circuits, whose complexity is almost independent of store size, is relatively unaffected.

It was therefore judged to be worthwhile to eliminate as much of the blanking intervals as the timing specification of the signals would allow.

BBC 625-line television transmissions are designed to meet the 'System I' specification⁶. According to this specification each line blanking interval has nominal duration of 12.05 μs (18.8% of the total line waveform). However, it is not possible to eliminate this interval entirely because considerable tolerances are allowed in the specification.

Fig. 2 shows an active television line with the extreme positions of the line blanking interval marked. By definition all times are referred to the mid-point of the leading edge of the line synchronising pulse (sync). From the figure it can be seen that the latest time at which line blanking may begin is $-1.3 \mu\text{s}$, and that the earliest it may

end is $+10.0 \mu\text{s}$. This leaves an interval of 11.3 μs that is guaranteed not to contain picture information.

A further inroad must be made because the signal is sampled, and the sampling sites do not in general coincide with the edges of the blanking interval. The sampling frequency proposed for the 34 Mbit/s system is twice colour-subcarrier frequency. Because this is not an exact multiple of line frequency, the sampling sites can occur almost anywhere on the line, and the whole sampling grid is tilted by about 0.15° from the vertical. Thus, only about 98 samples may be removed from each line blanking interval.

Another consequence of the non-integral number of samples per line is that it is not possible to retain and discard the same number of samples on every line. One or both of these numbers must be varied, and the choice may be influenced by the other processes being used to reduce the bit rate.

Originally it was thought that the number of samples discarded should be constant and divisible by two. This was because of constraints imposed at that time by the DPCM coder. Subsequently these constraints were removed, allowing the number of samples retained to be kept constant. This simplified the operation of the decoder.

The duration of the field blanking interval corresponds to 25 line periods plus one line blanking interval. Remembering that the line blanking intervals have already been accounted

for this represents a further 6.6% of the complete video waveform. However, because alternate field blanking intervals begin and end in the middle of active line periods, it was found simpler to retain these lines in their entirety, and to accept the small increase in bit rate (about 0.15%).

The final format is shown in Table 1. It repeats every four fields because this is the smallest number of fields to contain an exact number of samples. The number of samples deleted from the line blanking intervals alternates between 99 and 100. In practice the small increase from the permissible value of 98 is acceptable, because blanking

Natlock system), and may also be caused by monochrome signals that contain no colour burst for the sampling pulse generator to lock to. Modifications to the blanking format are dealt with in more detail in Section 3.1.2.

3. Description of the equipment

3.1. Coder

3.1.1. Hardware

A block diagram of the coder is shown in Fig. 3. It consists of a video buffer store, a con-

Table 1 — Pattern of samples removed from composite signal.

Field	Line	Samples Transmitted	Samples discarded from following blanking interval
1	23—309 inclusive	468	100 (even line numbers) 99 (odd line numbers)
1	310	468	14288
2	336—622 inclusive	468	99 (even line numbers) 100 (odd line numbers)
2	623	466	13724
3	23—309 inclusive	468	100 (even line numbers) 99 (odd line numbers)
3	310	468	14288
4	336—622 inclusive	468	99 (even line numbers) 100 (odd line numbers)
4	623	466	13723
Totals		539132	170247

intervals are usually longer than the minimum value, and because most receivers are set to over-scan. The total saving is 24.1% (17% line blanking, 7.1% field blanking). However, in a complete system some vestigial synchronising information would have to be reinserted to allow the blanking intervals to be correctly regenerated in the decoder. Thus, the net saving is estimated to be 23%.

The coder modifies the format as necessary to deal with changes in the ratio of line and sampling frequencies. This happens during certain source-synchronising operations (such as the BBC

troller, a synchronising-pulse detector and circuits to measure signal levels.

The buffer store smooths the flow of active video samples after the blanking intervals have been removed. Only active samples are written into the store; during the line and field blanking intervals the input to the store is disabled and the corresponding samples are discarded. At the same time samples are read from the store at a continuous but slower rate, chosen to be precisely equal to the long-term average input rate. In this way the occupancy of the buffer store is kept within

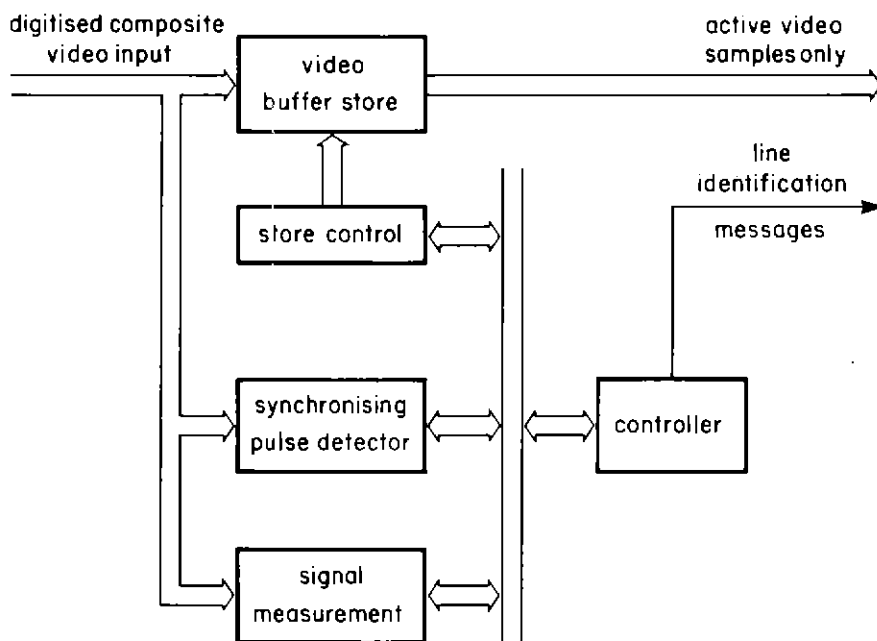


Fig. 3 — Main functions in coder.

bounds, and the size of the store could in principle be as small as 87 kbit (equivalent to about 19 complete lines of video). In practice a much larger store was used (262 kbit) partly to simplify the problems of simultaneous addressing for reading and writing, and partly in anticipation of other possible investigations using higher sampling rates.

It was decided at the outset of the work that a flexibility of design would be required that was difficult to achieve with conventional circuitry. For example, it would be necessary to develop strategies for non-ideal PAL signals and for recovery from switches between non-synchronous input signals. To avoid the possible need to rebuild parts of the equipment it was decided to put the main operating parameters under microprocessor control. Changes could then be made by changing the control program.

Ideally, the microprocessor would handle the video samples directly. However, even though the type chosen was one of the fastest available it was several hundred times too slow for this to be possible. Instead the microprocessor puts control information into interface registers that link it to the high-speed circuitry. Even with this reduced workload it was found that the programs had to be carefully written to achieve the required processes in the available time.

The video waveform is quantised to eight-bit accuracy, so that each sample is effectively converted into a number between 0 and 255. Line synchronising pulses (sync) are detected initially by searching for four successive sample values

below 32, which, in a video signal of the correct amplitude, is almost exactly midway between black and sync. levels. The need to detect four successive samples avoids spurious triggering on noise and large subcarrier excursions. Broad pulses and equalising pulses, where they begin in the middle of the line, are rejected by inhibiting the sync. detector until shortly before the line pulse is expected.

A similar detector, triggered at approximately the one-quarter and three-quarter points on the line, allows the controller to attain picture-lock by detecting the unique broad-pulse pattern on line 3.

The line-synchronising-pulse detector causes a counter to be reset every line. The counter is incremented at the video sampling rate and is used to count out the intervals to the centre of the synchronising pulse, to the front porch and to the colour burst so that the signal levels at these instants may be measured. If necessary the controller changes the threshold in the line-synchronising-pulse detector to maintain it mid-way between black and sync. levels.

The value that the counter has reached when it is reset is also recorded, to give the controller a line-by-line record of the number of sampling intervals elapsed. The controller uses this number to detect changes in line frequency relative to sampling frequency, so that it can modify the blanking format.

There is no direct link between the synchronising-pulse detector and the circuits

controlling access to the video buffer store. This is because the uncertainty in detecting each line synchronising pulse can be up to about two sampling intervals. It would be inefficient to modify the blanking format to track these random errors, because each of these changes would have to be signalled to the decoder. The coder therefore keeps to the basic blanking format unless a systematic drift is detected.

Because the stream of video samples has been stripped of all identification it is necessary to reinsert some information to enable the decoder to identify the individual active lines. In the experimental equipment this vestigial synchronising information was carried on a ninth wire, in parallel with the eight bits of video data. In a complete 34 Mbit/s coder it would be included in the data stream in the multiplexer and would be repeated several times to increase its immunity to errors. The synchronising information cannot be inserted in the video data stream in place of the original blanking samples because it would be corrupted by the DPCM coding and decoding.

3.1.2. The coder control program

The microprocessor chosen was the Signetics type 8 X 300⁷. Its operating speed of 250 ns per instruction allows 256 instructions to be executed per television line, and this has proved adequate.

Programs were written in MCCAP, an assembly specifically intended for the 8 X 300. Conversion from MCCAP to object code (cross-assembly) was performed using a purchased program run on a PDP11/23 mini-computer. MCCAP allows the use of 'macros': these are short sequences of instructions defined at the start of the program and given a name by which they can be called up in the program as often as required. This feature was used extensively to save labour in writing assembly code and was felt to be the best way of writing the program. A higher-level language tends to result in programs that are longer when they are turned into machine code. This is usually of no great concern but in this particular application, where speed is crucial, it would be important. Also, higher-level languages are not very suitable for writing control programs.

The flow chart of the coder's program is shown in Fig. 4. The program can be divided into two parts; the first part allows the equipment to lock to an incoming signal, the second part is executed once per television line once lock is achieved.

Lock is achieved in three stages, line lock, picture lock and eight-field lock.

To achieve line lock the program looks for eight lines of the correct duration. Eight lines are required to avoid the possibility of locking to broad pulses or equalising pulses in the field blanking interval.

To gain picture lock the program checks the broad-pulse detector for a single broad pulse in the first half of the line, which occurs only on line 3.

The 'System I' specification defines a numbered sequence of four fields in terms of colour-burst relative phase (normally 135° or 225°) and Bruch blanking (suppression of the colour burst on certain lines of the field blanking interval). The colour-burst phase is defined relative to a continuous (imaginary) subcarrier sinewave representing the U-axis. The four-field sequence can readily be identified by detecting the presence or absence of the colour burst on appropriate lines.

However, the four-field numbering sequence does not provide all the information needed by the decoder to regenerate the blanking interval waveform. The PAL waveform representing a stationary picture does not repeat every four fields, but every eight fields, so that two successive four-field 'System I' sequences are not identical. The difference between successive sequences is that all subcarrier components are reversed in phase, because four fields contain an odd number of subcarrier half-cycles. The 'System I' specification does not attempt to define any absolute phase relationship between colour subcarrier and the line or field synchronising pulses. It is therefore necessary to measure the phase of colour bursts with respect to, say, selected line synchronising pulses to extend the definition to eight fields.

An extra complication is that the colour bursts received by the coder are non-standard because they have been passed through a comb filter and sampled. (The comb filter removes frequencies from the video signal which cause aliasing when the signal is sampled at less than the Nyquist limit). The effect of the comb filter and the sampling is to remove the colour bursts on alternate lines and to increase the amplitude of those remaining by a factor $\sqrt{2}$. (A second comb filter in the receiving equipment restores the bursts to their original form.) This must be allowed for when measuring the amplitude of colour-burst and detecting its presence or absence for Bruch blanking. The phase reversal

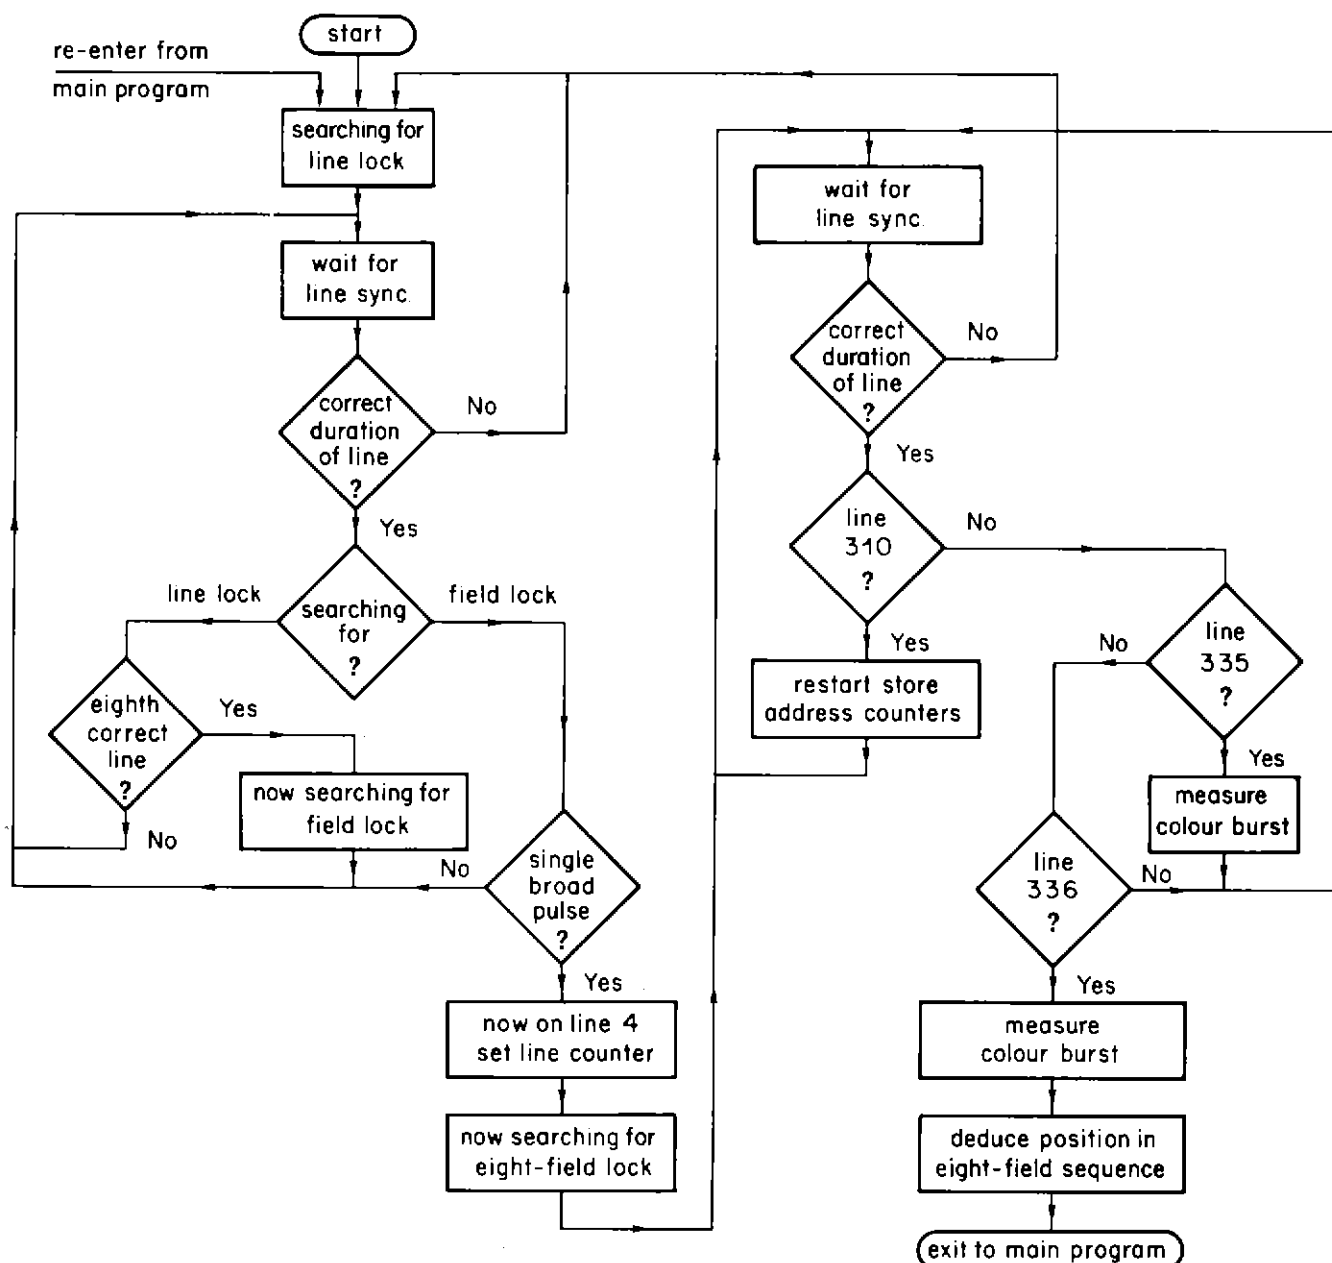


Fig. 4a – Coder: flowchart for attaining lock

of the remaining bursts between successive four-field sequences remains, however.

Once eight-field lock has been achieved the second part of the program is used, and this is executed once per television line unless the input signal is disturbed.

On most lines the program has four main tasks:

- (a) check the duration of the previous line;
- (b) prepare the circuits controlling access to the video buffer store for the next line;

(c) record sync. level, black level or colour burst amplitude;

(d) prepare an identifying header for the next line.

By keeping a tally of line durations the program checks that the line frequency is normal. Small discrepancies (an error of two sampling intervals or less) are ignored, but larger drifts cause the program to modify the blanking format. Otherwise the block of samples eliminated from the blanking intervals would begin to include active samples.

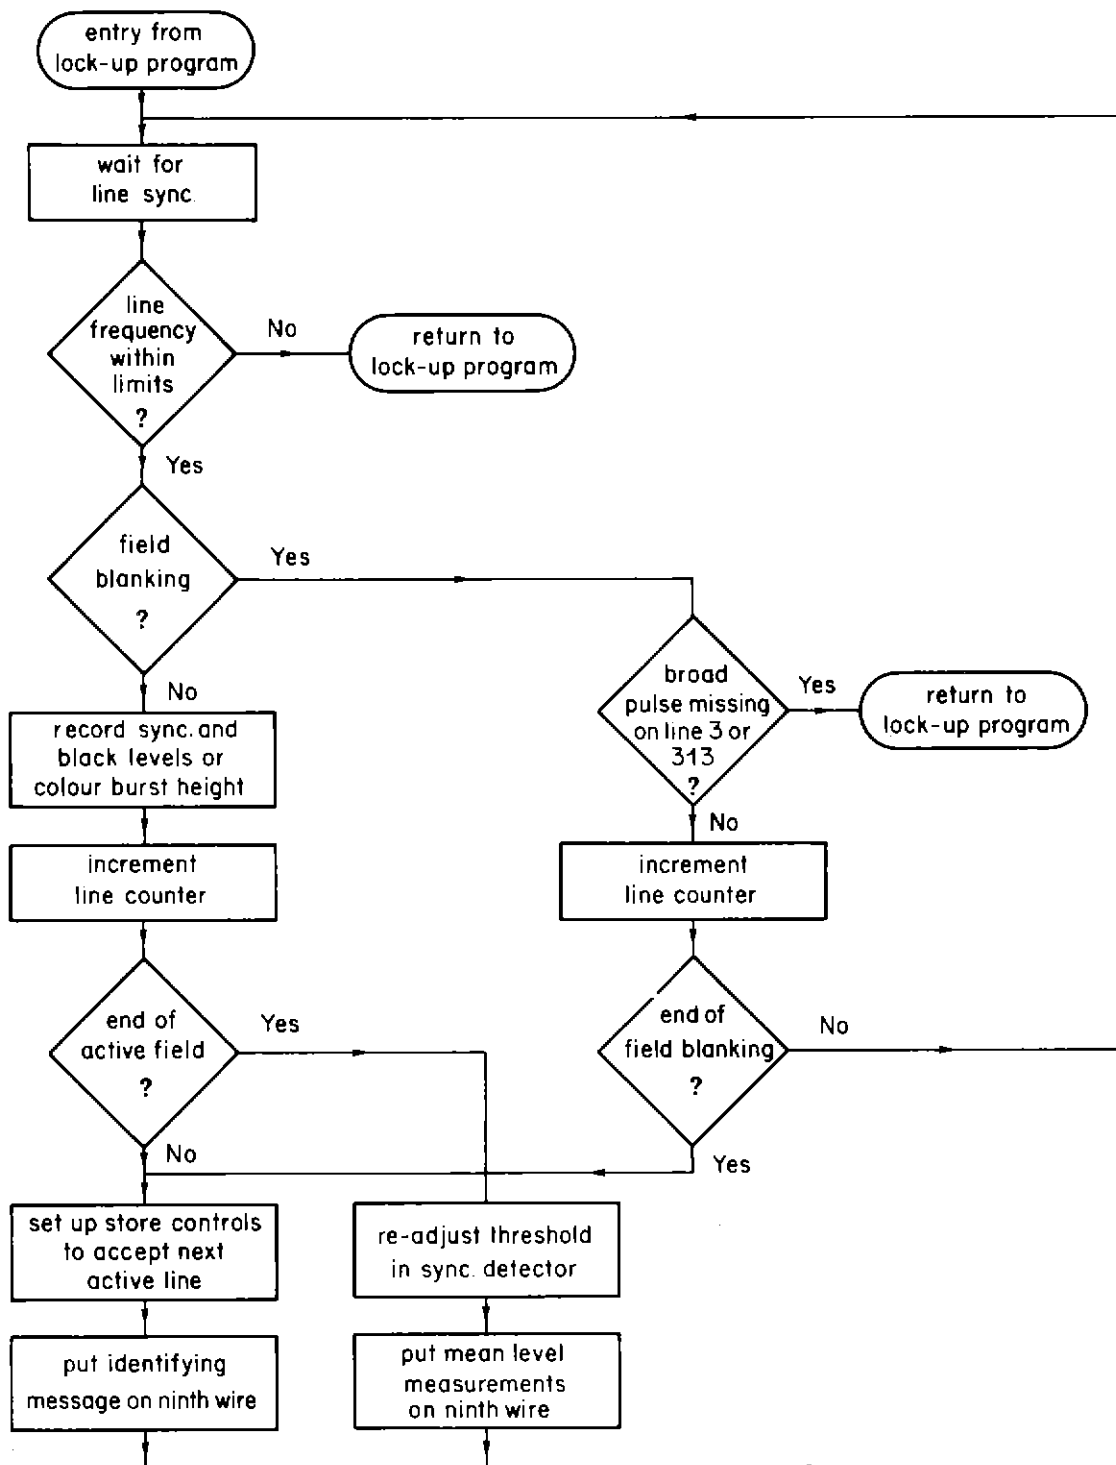


Fig. 4b — Coder: flowchart for main program.

Checking the line duration is made complicated because the true value (when expressed in multiples of the sampling interval) cannot be expressed exactly in binary notation. Even with 32-bit arithmetic an error would accumulate at a rate of about seven sampling intervals per hour. The technique adopted uses at most 16-bit arithmetic and has no long-term error. The correct

line duration corresponds to exactly 567.5032 sampling intervals, and on each line the controller subtracts either 567 or 568 from the tally. These numbers are subtracted in the proportions 621:629 respectively, so that over 1250 lines the average is $567\frac{629}{1250}$, which is exactly 567.5032. The remainder is always within one sampling interval of the correct value. The proportions 629:621

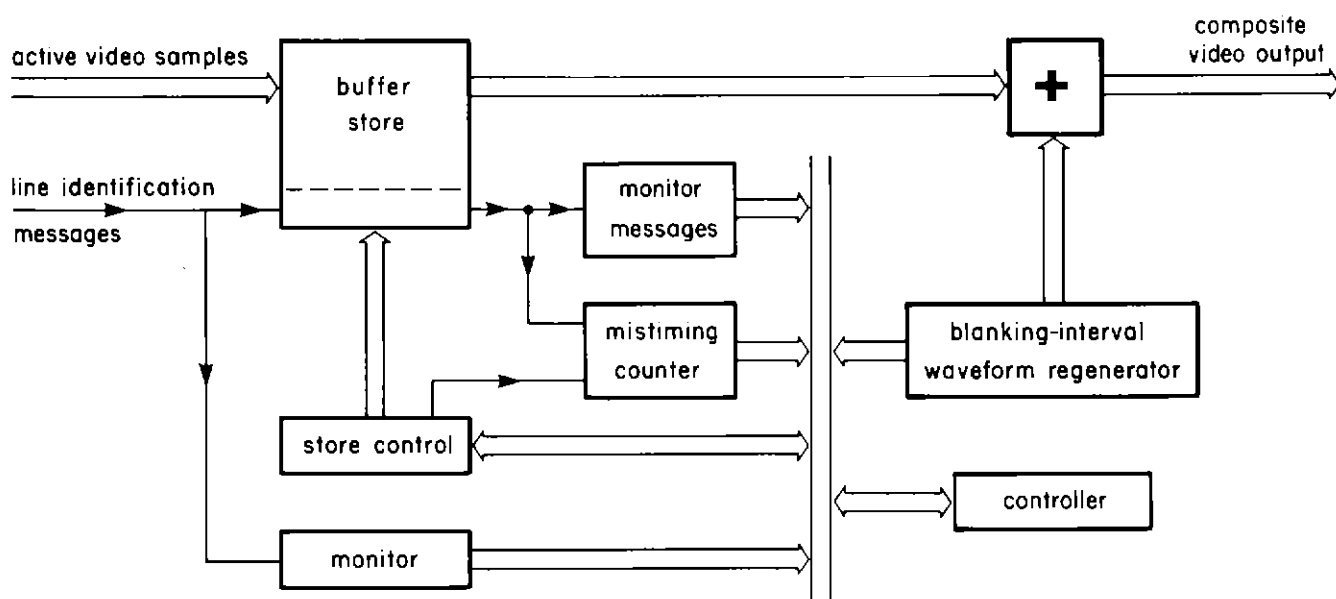


Fig. 5 — Main functions in decoder.

are derived by adding 629 every line to a separate register, modulo 1250. On those lines that a carry is generated 568 is subtracted; if no carry is generated, 567 is subtracted.

If a line synchronising pulse is found to be three sampling intervals (338 ns) or more from its expected position the blanking format is modified to track it. A disturbance of this size is unlikely to be caused by noise or by a change in synchronising pulse amplitude. The blanking format is modified by increasing or decreasing the number of samples discarded from the next line blanking interval. Because each modification causes a small change in the fraction of samples discarded the buffer store will ultimately overflow or be completely emptied. To prevent this, an adjustment is made to the number of samples retained on the last line of the field. The adjustment is calculated so that there is no long-term change in the ratio of samples discarded to samples retained.

If the line synchronising pulse is more than twenty sampling intervals from its expected position the controller assumes that the input signal has been disturbed. (The most likely cause is a switch between asynchronous sources). It therefore returns immediately to the first part of the program, to re-synchronise to the new signal. Similarly, the program is restarted if the single board pulses are not found on lines deemed to be numbers 3 and 313.

The measurements of colour burst amplitude, sync. and black level are each averaged over 128 lines per field. The combination of comb filtering

and sub-Nyquist sampling suppresses the colour burst on alternate lines, so the measurements of burst amplitude are alternated with those of sync. and black level. At the end of each field the threshold in the line-synchronising pulse detector is re-established midway between sync. and black level.

Each line written into the video buffer store is accompanied by an identifying message of 16 bits carried on a ninth wire. Four bits are reserved for internal 'housekeeping' and the remaining 12 bits are provided by the controller. On most lines the line and field number are transmitted. However, if the blanking format has been modified to track line-frequency drift this information is transmitted instead using reserved code words. Signal amplitudes and modifications to field blanking and the last active line of the field are transmitted in a block of lines near the end of the active field.

3.2. Decoder

3.2.1. Hardware

A block diagram of the decoder is shown in Fig. 5. Like the coder it contains a video buffer store and a microprocessor-based controller. The design of these components is identical to that of the coder. The decoder also contains circuits to detect the vestigial synchronising information transmitted by the coder and to regenerate the blanking interval waveforms.

The identifying messages are monitored at both the input and output of the buffer store.

Monitoring at the input is necessary to set up the addressing circuits of the store correctly. The messages are monitored at the output to ensure that complete lines are being read out. The line and field numbers are used to synchronise the circuits that regenerate the blanking interval waveform. When the decoder is correctly synchronised the last sample read from the buffer store before the blanking interval will be the last active sample of a line. A counter is provided which records any mistiming of the blanking interval. The control program uses this information when locking to a new input, and as a continuing check on its own operation.

The circuits that regenerate the synchronising pulses and the colour burst are shown in Fig. 6. These parts of the waveform are basically repetitive and well-defined, but they are allowed to vary within limits set by the System I specification. It is important that any variation on the input signal to the coder is faithfully reproduced at the decoder. Otherwise circuits that use burst and sync. amplitude, for example, to equalise chrominance and luminance gain will not operate correctly.

An analogue link can accommodate reasonable amplitude and frequency errors in the signal it is carrying. If the signal amplitude falls the signal-to-noise ratio will deteriorate. If it rises the link will begin to clip the signal momentarily on switching between dark and light pictures. However, the overload margin is high, usually 3 dB. The variations in the line frequency used for

source synchronising are small (about 1 part in 5600) and are not affected by link characteristics.

It is therefore important that a digital link can also tolerate such variations in input signal parameters with minimal distortion of the signal. It is also important that any degradation is gradual and that the link does not suddenly fail as a parameter drifts out of tolerance.

The regenerator was therefore designed to be programmed to produce synchronising pulses and colour bursts of any amplitude, subject only to the digital coding range and to the limits of eight-bit quantising. Line frequency could be varied from 1/568 to 1/567 of the sampling frequency (a variation of $\pm 0.09\%$ of its nominal value). For simplicity the regenerator was programmed to operate only at the nominal line frequency, deviations being accommodated by occasional phase shifts between the active and regenerated portions of the signal.

Because the sampling frequency is not a simple multiple of the nominal line frequency a different sequence of samples must be generated for every line. (A line-locked sampling frequency would not avoid this problem. It would merely transfer it to the regeneration of the subcarrier sine wave in the colour burst.)

The method of generating the synchronising pulse edges has been described elsewhere⁸. The edge of a synchronising pulse is an integrated raised cosine, that is, $x - \sin x$, and a straight-line

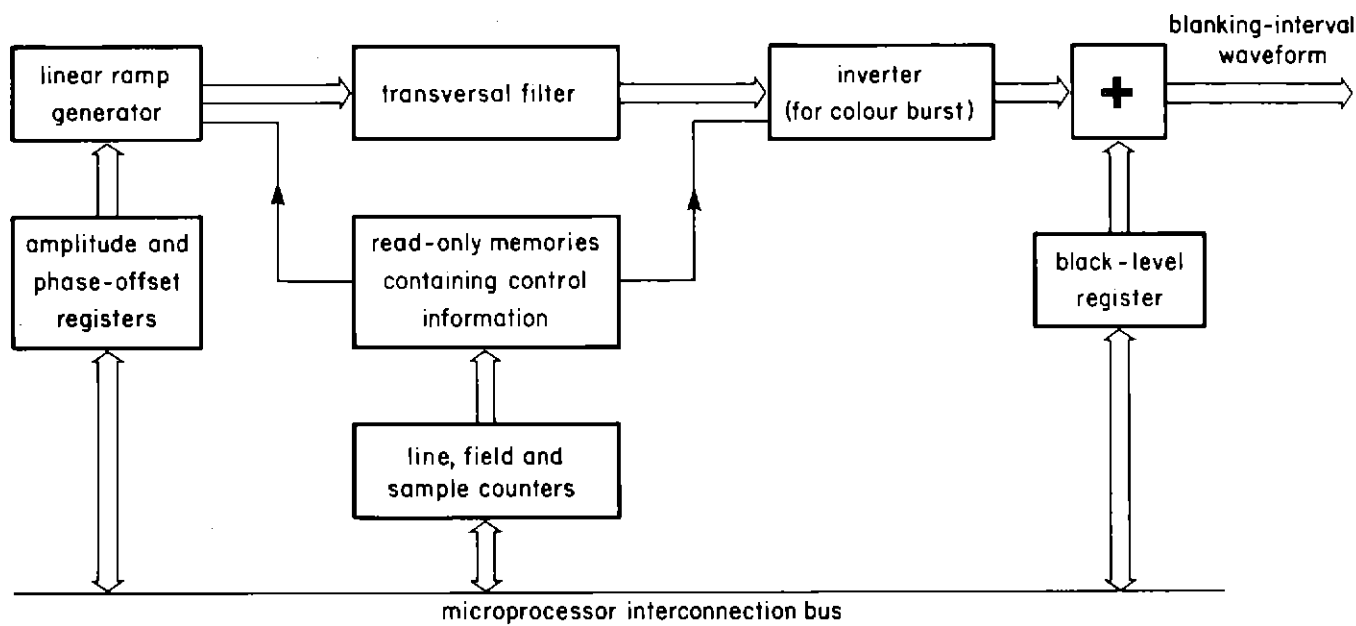


Fig. 6 — Regeneration of blanking interval waveform.

approximation is inadequate. The method adopted is to apply a three-step ramp to a three-term transversal filter. Timing of the edge relative to the sampling intervals is adjusted by adding an offset to the first two steps of the ramp, this offset being calculated line-by-line by the controller.

The colour burst is regenerated first as a pulse, effectively its envelope, to which subcarrier modulation is applied by inverting every other $2f_{sc}$ sample about black level.

The line and field numbers are stored in registers, and details of the blanking interval of each line are stored in read-only memory. These details are the presence or absence of colour burst and the position of any broad pulses or equalising pulses. The line- and field-number registers can

be loaded by the controller to synchronise pulse generation to the video waveform. The controller can also read from the registers to re-check synchronisation.

3.2.2. The decoder control program

The flow chart of the control program for the decoder is shown in Fig. 7. Like the control program for the coder it is divided into two parts, initial synchronisation and routine operations.

First, the controller checks that the decoder is receiving a proper input signal by waiting for and comparing successive line messages. The line numbers should, of course, be consecutive. Then, it searches for a line shortly after the start of the active field interval, a convenient point at which to

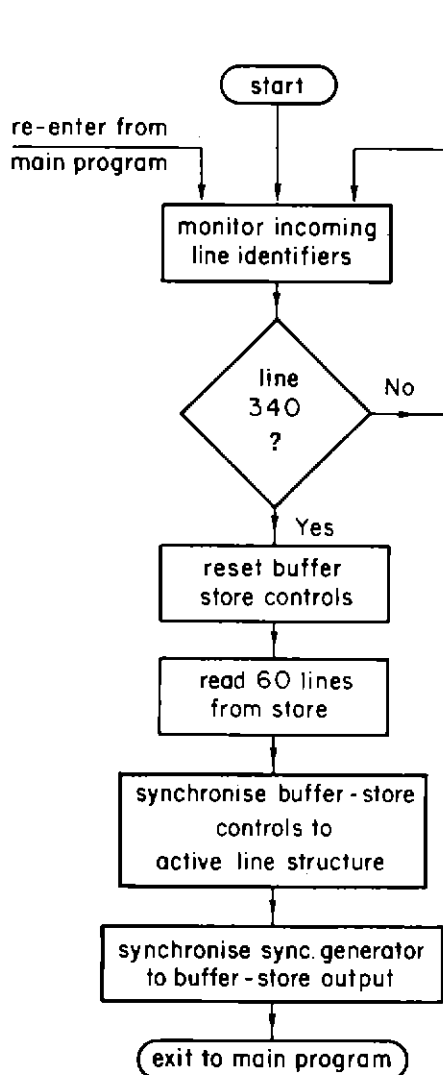


Fig. 7a – Decoder: flow-chart for attaining lock.

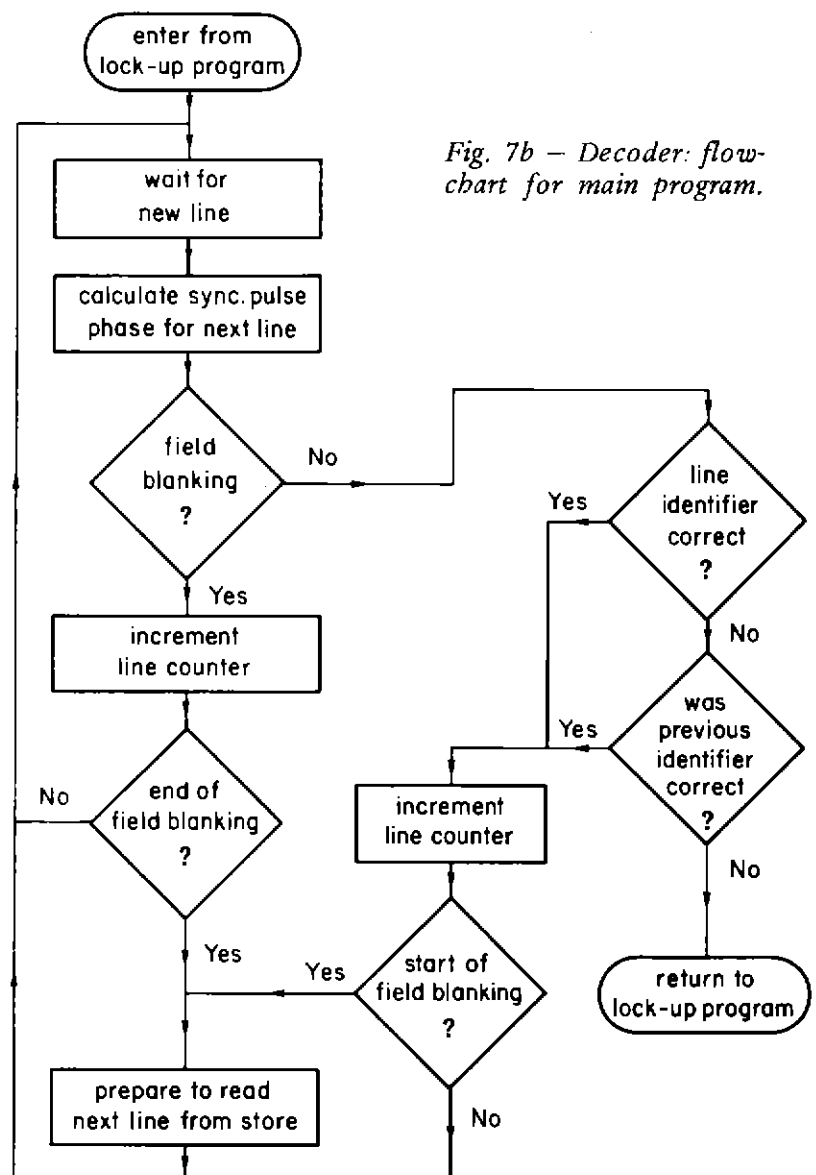


Fig. 7b – Decoder: flow-chart for main program.

reset the address registers for the video buffer store. The occupancy of the store should be at its highest at this point, and the read-address will be set as far behind the write-address as possible. Otherwise the store will be emptied before the end of the active field. The controller then reads sixty lines from the store to clear the store of any samples written into it before the address generators were correctly set.

The moment at which the address registers were reset relative to the video line timing is not defined. This is because the time taken to reach that part of the control program will vary. When the controller reads the first line from the store it will probably consist of the last part of one line followed by the start of another. This will be recorded by the mis-timing counter, and the controller will adjust the length of the next read-cycle to synchronise the active and blanking intervals.


The line and field numbers in the identifying message are transferred to the blanking-interval waveform regenerator. This completes the initial setting up, and the following routines are executed once per line until the input signal is disturbed:

- (a) a check that the identifying messages read from the buffer store are consistent;
- (b) a check that the last sample read from the store on each line is actually the last active sample of that line;
- (c) preparation of the circuits controlling the buffer store to read out the next line at the correct time;
- (d) an update of the phase offsets in the blanking interval waveform regenerator.


4. Performance

Fig. 8 shows test-card pictures artificially displaced to show the blanking intervals. Fig. 8(a) shows the normal picture signal. In Fig. 8(b) the blanking intervals have been suppressed but the monitor is still synchronised to the normal signal. In Fig. 8(c) the blanking is suppressed but the monitor has been modified to run at the lower line frequency to display a recognisable picture.


The following measurements were made with the coder and decoder connected back-to-back. Analogue/digital conversion was by a BBC Research Department type EP1M/535 codec⁹,



(a)



(b)




(c)

Fig. 8 — Test-card pictures artificially displaced to show blanking intervals.

- (a) Normal picture signal.
- (b) Picture signal with blanking suppressed ('short' video).
- (c) Picture signal with blanking suppressed and viewed on a specially-adapted television monitor.

Fig. 9 — Effect of noise on analogue input signal.



sampling at twice subcarrier frequency. Analogue comb filters were used.

4.1. Errors in input signal level

The largest analogue signal that can be accepted was limited by the overload margin of the analogue-to-digital converter. This margin is kept as small as possible, to leave the largest number of quantising levels for the normal signal range, thus minimising quantising noise. The coder was set to overload at approximately ± 0.4 dB relative to the standard video signal amplitude of 700 mV between black and white levels. The smallest signal to which the codec would maintain or attain lock was -1.2 dB.

4.2. Noise on analogue input

White noise over the frequency band 30 Hz to 5.5 MHz was added to the analogue video input signal. The drop-out rate was measured for a range of noise levels.

The results of these measurements are shown in Fig. 9. The effect of noise is to interfere with the detection of line synchronising pulses and broad pulses. As the noise power is reduced the probability of the instantaneous noise voltage causing the signal to cross the sync. threshold falls, and this is seen in the results.

However, no fall was observed beyond signal-to-noise ratio (peak signal to r.m.s. unified weighted noise) of about 40 dB. The reason for this was not discovered during the time available for the tests (at this noise level the decoder was losing lock only about once every three hours). A mains filter was fitted but this had no observable effect. Over such long periods of time it was, of course, difficult to ensure that the input signal itself was totally undisturbed.

4.3. Variations in line and field frequency

Variations in line and field frequency were simulated by altering the sampling frequency of the analogue-to-digital converter. This is not an exact simulation because the sampling frequency is no longer locked to the colour subcarrier frequency. However, this does not effect those parts of the circuitry and program concerned with locking to the synchronising pulses.

A variation of +1 part in 1062 and -1 part in 1083 could be tolerated. The normal 'fast' rate for Natlock is approximately ± 1 part in 5600¹⁰.

4.4. Non-synchronous 'cuts'

The ability of the coder and decoder to regain lock after a disturbance was tested by

switching between two asynchronous sources. The lock-up time was measured by observing the analogue input and output signals on a storage oscilloscope, at a slow sweep speed.

The lock-up time varied according to the relative phasing of the two signals. The longest lock-up time observed was approximately 150 ms, and on most cuts the equipment regained lock within 80 ms.

4.5. Propagation delay

The coder and decoder both contain large buffer stores to redistribute the active video samples. These stores are responsible for most of the delay in the video path. The delay through the coder and decoder connected back-to-back was measured as $5.2 \text{ ms} \pm 0.1 \text{ ms}$. This figure includes the delay through the comb pre- and post-filters, which is about $64 \mu\text{s}$.

5. Conclusions

A reduction of 23% can be made in bit rate by eliminating the line and field blanking intervals from a digitised composite-PAL video signal; a reduction of 17% can be made by eliminating the line blanking intervals only. Both figures take into account timing tolerances and include an allowance for vestigial synchronising information for regenerating the missing parts of the waveform.

For a target bit rate of 34 Mbit/s for a video signal and ancillaries it was judged worthwhile to make the larger reduction. The main penalty was extra storage, which was relatively easy to provide. The advantage was that the available number of bits per sample could be increased from 4.2 to 4.5. At the level of data compression required even this small increase would significantly improve the picture quality or, alternatively, considerably reduce the complexity of the DPCM predictor.

The equipment built for the study has been described and its performance discussed. Micro-processor control has been used, partly to allow changes to be made during development and partly because of the complexity of operation of the equipment. The complexity is necessary to accommodate tolerances on the analogue signal which must be deliberately allowed for in digital processing, although they would be acceptable in analogue transmission. Further complication is caused by the non-integral relationship between line frequency and the sampling frequency, and the absence of a defined phase relationship between them.

This complexity makes the elimination of blanking intervals somewhat unattractive as a means of reducing the bit rate of a digitised composite-PAL signal. However, without blanking elimination it would be very difficult indeed to compress a video signal to 32 Mbit/s without affecting the picture quality. The problems could largely be overcome by using a digital field-store synchroniser to 'standardise' the input signal to conform to the System I specification. However, this merely transfers complexity from one area to another, would probably increase the overall cost and may have operational disadvantages.

On digital signals that are more tightly defined, for example those complying with the CCIR recommendations for the digital coding of component signals, blanking elimination would be considerably simplified and is already being considered.

6. References

1. DEVEREUX, V.G. 1971 Pulse-code modulation of video signals: subjective study of coding parameters. BBC Research Department Report No. 1971/40.
2. DEVEREUX, V.G. and PHILLIPS, G.J. 1974. Bit-rate reduction of video signals using differential p.c.m. techniques. *IBC '74, IEE Conference Publication No. 119*, pp 83-89.
3. STOTT, J.H. and RATLIFF, P.A. 1983. Digital television transmission: 34 Mbit/s PAL investigation. BBC Research Department Report in preparation.
4. WELLS, N.D. 1983. Digital video bit-rate reduction: digital comb filters for sub-Nyquist ($2f_{sc}$) PAL. BBC Research Department Report in preparation.
5. STOTT, J.H. 1983. Digital video bit-rate reduction: DPCM for sub-Nyquist ($2f_{sc}$) PAL. BBC Research Department Report in preparation.
6. 1971. Specification of television standards for 625-line System I transmissions. Joint BBC/ITA publication.
7. NEMEC, J. 1977. One-chip bipolar micro-controller approaches bit-slice performance. *Electronics International*, Vol. 50, No. 18 pp 91-96.

8. CROLL, M.G. 1978. Digital video: bit-rate reduction by removal of the line blanking portion of the waveform. BBC Research Department Report No. 1978/17.
9. DEVEREUX, V.G. 1970. Pulse-code modulation of video signals: 8-bit coder and decoder. BBC Research Department Report No. 1970/25.
10. BLISS, J.L., MILLAR, I.D.B., ALLEN, C.P. 1968. Picture source synchronisation; the NATLOCK system. EBU Review No. 107A, February 1968, pp 2 to 7.

